



AP
ICW
\$

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

PATENT APPLICATION

ATTORNEY DOCKET NO. 10003533-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor(s): Lance W. Russell

Confirmation No.: 9456

Application No.: 09/888,544

Examiner: Barqadle, Yasin M.

Filing Date: June 25, 2001

Group Art Unit: 2153

Title: Routing Meta Data for Network File Access

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Transmitted herewith is the Appeal Brief in this application with respect to the Notice of Appeal filed on Feb. 28, 2006 and the Notice of Panel Decision from Pre-Appeal Brief Review dated April 11, 2006. The fee for filing this Appeal Brief is (37 CFR 1.17(c)) \$500.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

☐ (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d)) for the total number of months checked below:

☐ 1st Month
\$120

☐ 2nd Month
\$450

☐ 3rd Month
\$1020

☐ 4th Month
\$1590

☐ The extension fee has already been filed in this application.

☒ (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account 08-2025 the sum of \$ 500. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

☒ I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to:
Commissioner for Patents, Alexandria, VA 22313-1450
Date of Deposit: May 5, 2006

OR

☐ I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number (571)273-8300.

Date of facsimile:

Typed Name: Edouard Garcia

Signature: [Signature]

Respectfully submitted,

Lance W. Russell

By [Signature]

Edouard Garcia

Attorney/Agent for Applicant(s)

Reg No. : 38,461

Date : May 5, 2006

Telephone : (650) 289-0904



UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Lance W. Russell
Serial No. : 09/888,544
Filed : June 25, 2001
Title : Routing Meta Data for Network File Access

Art Unit : 2153
Examiner : Barqadle, Yasin M.

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

I. Real Party in Interest

The real party in interest is Hewlett-Packard Development Company, L.P., a Texas Limited Partnership having its principal place of business in Houston, Texas.

II. Related Appeals and Interferences

Appellant is not aware of any related appeals or interferences that will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

III. Status of Claims

Claims 1, 2, 4-6, 12, 13, 19, and 21-27 are pending.

Appellant appeals all rejections of the claims 1, 2, 4-6, 12, 13, 19, and 21-27.

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as First Class Mail in an envelope addressed to: Commission for Patents, PO Box 1450, Alexandria, VA 22313-1450 on:

May 5, 2006

Date

(Signature of person mailing papers)

Edouard Garcia

(Typed or printed name of person mailing papers)

05/10/2006 09888544 082025 500.00 DA
01 FC:1402

IV. Status of Amendments

The Amendment that was filed on September 6, 2005, has been entered and acted upon by the Examiner.

No amendments were filed after the final Office action dated December 2, 2005.

V. Summary of Claimed Subject Matter

The invention claimed in independent claim 1 is a method of accessing a data file in a distributed computing environment. In response to a request from a client site for access to a data file stored in one or more physical storage systems at a source site, physical address meta data and routing meta data are sent from the source site to the client site. The physical address meta data includes physical addresses of one or more logical blocks of the data file in the one or more physical storage systems. The routing meta data includes one or more node addresses along one or more network routes between the client site and the source site. The routing meta data allows the client site to optimize the selection of routes over which the data file is accessed based upon client-specific criteria (e.g., packet delay and fragmentation criteria). In this way, embodiments in accordance with the claimed invention can avoid the sub-optimal route selection that may occur with table-based network routing protocols in which routing decisions are made based upon network traffic considerations, rather than file system considerations. In addition, embodiments in accordance with the claimed invention can avoid the network overhead (e.g., file system access to routing tables in addition to normal network traffic access) that otherwise would be required to select optimal routes using such table-based network routing protocols.

FIG. 1 shows a distributed computing system 10 that includes a client site 12 that is connected to a source site 14 through two intermediate networks 16, 18 (see page 5, line 25 - page 6, line 28 of the specification). Each of the client site 12 and the source site 14 includes a respective client system 20, 22, a respective file system 24, 26, a respective meta data system 28, 30, a respective block server 32, 34, and one or more respective physical storage systems 36, 38. In the illustrated embodiment, the physical storage systems 38 of the source site 14 contain one or more data files that are accessible by the client site 12 (see page 6, line 25 - page 7, line 1 of the specification).

FIG. 3 shows an embodiment of a method by which the source site 14 handles file access requests from the client site 12 in accordance with the invention defined in independent claim 1. In response to a request from the client site 12 for access to a data file stored in the physical storage systems 38 at the source site 14, the source file system 26 passes the file request to the source meta data system 30 (FIG. 3, block 82). The source meta data system 30 passes to the source file system 26 meta data that is associated with the requested data file (FIG. 3, block 84; see page 9, lines 4-7 of the specification). This meta data includes access protection meta data, physical address meta data, and routing meta data. After confirming that the client site 12 is authorized to access the requested data file (FIG. 3, blocks 86-88), the source file system 26 sends to the client site 12 a reply containing meta data associated with the logical file blocks of the requested data file (FIG. 3, block 92; see page 9, lines 12-15 of the specification). This meta data includes the physical address meta data and the routing meta data.

With reference to FIG. 2, "the client file system 24 selects an optimal network route to source site 14 based upon routing meta data associated with the physical address information for the requested data file (step 62)" (page 7, line 29 - page 8, line 1 of the specification). The client file system 24 may select a network route to source site 14 based upon packet delay and fragmentation criteria, such as the load conditions, transmission characteristics, and MTU sizes of intermediate networks 16, 18 (see page 8, lines 1-6 of the specification). "Therefore, even if the file system uses the same routing algorithms as the general networking code to determine optimal routes, the file system may select a different path than the general network code" (page 8, lines 6-9). "After selecting an optimal network route over which to access the requested data file (step 62), client file server 24 accesses the logical file blocks for the requested data file through source block server 34 and returns the logical file blocks to the client application program (step 64)" (page 8, lines 9-13).

VI. Grounds of Rejection to be Reviewed on Appeal

A. Claims 1, 2, 6, 12, 13, and 19 stand rejected under 35 U.S.C. § 102(e) over Vahalia (U.S. 2005/0251500).

B. Claims 4, 5, 23, 24, 26, and 27 stand rejected under 35 U.S.C. § 103(a) over Vahalia in view of Koyanagi (U.S. 20010013067).

C. Claims 21, 22, and 25 stand rejected under 35 U.S.C. § 103(a) over Vahalia in view of Kato (U.S. 6,223,249).

VII. Argument

A. Rejection under 35 U.S.C. § 102(e) over Vahalia (U.S. 2005/0251500)

The Examiner has rejected claims 1, 2, 6, 12, 13, and 19 under 35 U.S.C. § 102(e) over Vahalia (U.S. 2005/0251500).

1. Overview of Vahalia's disclosure

Vahalia discloses a network file server architecture that allows shared data access to files stored in a file system. Vahalia describes his invention in the context of a prior art file server 20, which is shown in FIG. 1. The file server 20 allows each client 26, 27 to access the same read/write file in a file system 24 through two data movers 21, 22. Each data mover 21, 22 performs file locking management and mapping of the network file to logical blocks of storage in the cached disk storage 25 and moves data between the client and storage (see ¶ 6). Each data mover 21, 22 provides at least one network port for servicing client requests (see ¶ 7). "The cached disk storage 25 is configured so that the file system 23 is accessible only through the data port connected to the first data mover 21 and so that the file system 24 is accessible only through the data port connected to the second data mover 22" (¶ 9). In the prior art file server 20, the clients 26, 27 communicate with the data movers 21, 22 using the NFS protocol (see ¶ 12).

For illustrative purposes assume that the file system 23 contains a give read/write file. The data mover 21, which owns (i.e., has exclusive access to) the file system 23, provides the client 26 with access to the given read/write file by placing a lock on the file, accessing the file in the file system 23, and streaming the read/write data to the client 26 (see ¶ 10). The data mover 22, on the other hand, provides the client 27 with access to the given read/write file by acting as a proxy router that forwards NFS data packets from the client 27 to the data mover 21 (see ¶ 48).

By “using a data bypass path around the data mover that owns the file system during transmission of read/write data,” Vahalia’s invention provides significant improvement in data access time compared to the prior art file server 20 (§ 50). In accordance with Vahalia’s teachings, the data bypass path is implemented by a “high-speed data link,” which is a term that Vahalia consistently uses to refer to a dedicated, point-to-point link that is designed for the high-speed exchange of read/write data (see, e.g., § 11, where Vahalia refers to the dedicated, point-to-point read/write data link between the data movers 21, 22 as a “high-speed data link”).

It was well-known in the art of network communications at the time the invention was made that physical addresses conventionally are not used when transmitting frames over point-to-point communication links because there is only one possible destination for each transmission. In addition, host addresses conventionally are not assigned to the nodes of a point-to-point communications link. Although IP routing tables may use arbitrary values as next-hop addresses for the nodes of a point-to-point communications link, such values are ignored by both IP and the point-to-point hardware interfaces of these nodes. Vahalia does not teach or suggest anything that contradicts this well known information about point-to-point communications links.

2. Independent claim 1

Claim 1 recites:

1. A method of accessing a data file in a distributed computing environment, comprising:

in response to a request from a client site for access to a data file stored in one or more physical storage systems at a source site, sending from the source site to the client site physical address meta data including physical addresses of one or more logical blocks of the data file in the one or more physical storage systems, and routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site.

The Examiner’s rejection of claim 1 under 35 U.S.C. § 102(e) over Vahalia should be withdrawn because Vahalia does not disclose a method that includes sending from a source site to a client site routing meta data comprising one or more node addresses along one or

more network routes between the client site and the source site in response to a request from the client site for access to a data file stored in one or more physical storage systems at the source site.

3. The Examiner's position and Appellant's preliminary rebuttal

a. Introduction

The Examiner has taken the position that each and every element of independent claim 1 is disclosed by Vahalia in ¶¶ 15 and 81-89 (see page 3, lines 7-20 of the final Office action in which the Examiner has quoted ¶ 15 in its entirety). Contrary to the Examiner's position, however, none of the cited paragraphs teaches the subject matter defined in claim 1.

b. Vahalia, ¶ 15

In ¶ 15, Vahalia discloses that a client is connected to a server by an IP data link and is additionally connected to data storage by a high-speed data link that bypasses the server. In accordance with the method described in ¶ 15, the client uses a file access protocol to obtain from the server metadata specifying the data storage locations of a requested data file. The client produces a data access command from the metadata obtained from the server. The client then sends the data access command to the data storage over the high-speed data link using a high-speed data protocol.

On its face, ¶ 15 does not disclose each and every feature of independent claim 1. In particular, the metadata that is sent to the client by the server only corresponds to the physical addresses of the logical blocks of the requested data file; this metadata does not include "routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site," as recited in claim 1. Indeed, Vahalia explains that (¶ 52):

The term metadata refers to information about the data, and the term metadata is inclusive of file access information and file attributes. The file access information includes the locks upon the files or blocks of data in the files. The file attributes include pointers to where the data is stored in the cached disk array.

Thus, the metadata sent by Vahalia's server does not include routing metadata comprising one or more node addresses along one or more network routes between the client and the server. Moreover, the inclusion of such routing metadata in the transmission from the server in response to the client's file access request would not serve any purpose whatsoever in the context of Vahalia's file server system because physical addresses and IP addresses are ignored in communications over point-to-point, high-speed data links in accordance with the knowledge that was generally available at the time the invention was made.

c. Vahalia, ¶¶ 81-89

¶¶ 81-89 also do not disclose sending from a source site to a client site routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to a request from the client site for access to a data file stored in one or more physical storage systems at the source site. Instead, ¶¶ 81-89 merely describe a process of using the CIFS protocol for sharing data sets among data movers (see ¶ 68).

i. Vahalia, ¶ 81

¶ 81 describes the general process by which a client accesses a file on a server in accordance with the CIFS protocol. In this process, (¶ 81):

... a client has to: (1) parse the full file name to determine the server name, and the relative name within that server; (2) resolve the server name to a transport address (this may be cached); (3) make a connection to the server (if no connection is already available); and (4) exchange CIFS messages.

None of the steps (1)-(4) of this process involves transmitting from the server to the client routing meta data comprising one or more node addresses along one or more network routes between the client and the server.

In accordance with the CIFS protocol, the server name resolving step (1) merely involves having the client parse a URL for the server name and the relative name. In one commonly cited example of server name parsing in accordance with the CIFS protocol, the URL "file://fs.megacorp.com/users/fred/stuff.txt" is parsed by the client into the server name

“fs.megacorp.com” and the relative name “/users/fred/stuff.txt” based on cues from the double slashes and the next slash.

The server name resolving step (2) merely involves resolving of the server name to a transport address. In accordance with the CIFS protocol, this step may be performed using any name resolution mechanisms, such as the DNS. As was well known in the art at the time the invention was made, the process of resolving a server name into a transport address does not involve transmitting from the server to the client routing meta data comprising one or more node addresses along one or more network routes between the client and the server. Instead, such a service is provided by a name resolver server (e.g., a DNS server). In addition, the transport address of the server is not a node address along one or more network routes between the client and the server.

The connection making step (3) merely involves executing a prescribed handshaking procedure between the client and the server. In ¶ 89, Vahalia teaches that the CIFS protocol establishes connections using the NETBIOS session service. As was well known in the art at the time the invention was made, this process of establishing a connection between the client and the server does not involve transmitting from the server to the client routing meta data comprising one or more node addresses along one or more network routes between the client and the server.

The CIFS message exchanging step (4) merely involves transmitting data between the client and server during an established session. In accordance with the NETBIOS session protocol, each data packet is responded to with either an acknowledgement packet or a negative acknowledgement packet. As was well known in the art at the time the invention was made, the process of exchanging CIFS messages does not involve transmitting from the server to the client routing meta data comprising one or more node addresses along one or more network routes between the client and the server.

ii. Vahalia, ¶¶ 82-88

¶¶ 82-88 describes the format of server message blocks (SMBs) that the client exchanges with the server in accordance with the CIFS protocol.

iii. Vahalia, ¶ 89

¶ 89 describes the initial step in a CIFS request sequence for request forwarding.

With reference to FIG. 7, ¶ 89 explains that:

In a first step 131, in response to a file access request from a client, the network opens a TCP connection between the client and the server. As described in Leach, Appendix A, p. 119-120, this includes resolving the server name in the client request to an IP address of the Forwarder, and establishing a connection between the client and the Forwarder if a connection has not already been set up. Connection establishment is done using the NETBIOS session service, which requires the client to provide a "calling name" and a "called name."

As explained above, the process of resolving the server name merely involves resolving of the server name to a transport (IP) address. In accordance with the CIFS protocol, this process may be performed using any name resolution mechanisms, such as the DNS. As was well known in the art at the time the invention was made, the process of resolving a server name into an IP address does not involve transmitting from the server to the client routing meta data comprising one or more node addresses along one or more network routes between the client and the server. Instead, such a service is provided by a name resolver server (e.g., a DNS server). In addition, the IP address of the server is not a node address along one or more network routes between the client and the server.

Similarly, the process of establishing a connection between the client and the server merely involves executing a prescribed handshaking procedure in accordance with the NETBIOS session service. As was well known in the art at the time the invention was made, this process of establishing a connection between the client and the server does not involve transmitting from the server to the client routing meta data comprising one or more node addresses along one or more network routes between the client and the server.

d. Conclusion

Thus, none of the paragraphs that are cited by the Examiner supports his position that Vahalia discloses each and every feature of the subject matter recited claim 1. Indeed, as

explained in detail below, there is no teaching of the subject matter defined in claim 1 anywhere in Vahalia's disclosure.

4. Detailed analysis of Vahalia's disclosure

a. Introduction

Vahalia discloses several embodiments of a network file server that allows a client to access a data file stored in a file system of a cached disk array. The different file server embodiments that are disclosed in Vahalia are shown in FIGS. 1-4. None of these embodiments, however, sends to the client site routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to a request from a client site for access to a data file stored in one or more physical storage systems at the source site, as recited in claim 1. Consequently, none of Vahalia's embodiments provides the advantages that are enabled by the invention defined by claim 1, such as the avoidance of sub-optimal route selection that otherwise may occur with table-based network routing protocols and the avoidance of unnecessary network overhead that otherwise would be required using table-based network routing protocols.

b. The embodiment disclosed in FIG. 1 of Vahalia

In the prior art embodiment shown in FIG. 1, a client 26 connects to a data mover 21 over a data network 30. In response to a request from the client 26 for access to a data file, the data mover 21 either accesses the data file directly from file system 23 or indirectly from file system 24 through the data mover 22, which owns the file system 24. In both of these cases, the data mover 21 streams read/write data between the client 26 and the file system in which the data file is stored over the existing network connection (see ¶¶ 10, 48, 49, 95).

In this process, the data mover 21 does not send to the client 26 routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to the request to access the data file. Indeed, the transmission of such routing metadata from the data mover 21 in response to the client's file access request would not serve any purpose whatsoever in the context of the embodiment

shown in FIG. 1 because a session already has been established between the client 26 and the data mover 21 as described above in § VII.3.c, and the client 26 does not establish another network connection over another network route after the initial connection with the data mover 21 has been established.

c. The embodiment disclosed in FIG. 2 of Vahalia

In the embodiment shown in FIG. 2, a client 46 connects to a data mover 41 over a data network 50. In response to a request from the client 46 for access to a data file, the data mover 41 accesses the data file either directly from file system 43 or directly from file system 44 over a high-speed direct (point-to-point) data bypass path 48 that bypasses the data mover 42, which owns the file system 44. In both of these cases, the data mover 41 streams read/write data between the client 46 and the target file system in which the data file is stored over the existing network connections between the client and the target file system (see ¶ 53).

In this process, the data mover 41 does not send to the client 46 routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to the request to access the data file. Indeed, the transmission of such routing metadata from the data mover 21 in response to the client's file access request would not serve any purpose whatsoever in the context of the embodiment shown in FIG. 2 because a session already has been established between the client 26 and the data mover 21 as described above in § VII.3.c, and the client 46 does not establish another network connection over another network route after the initial connection with the data mover 41 has been established.

d. The embodiment disclosed in FIG. 3 of Vahalia

In the embodiment shown in FIG. 3, a client 64 connects to a data mover 61 over a data network 70 and connects to a file system 62 owned by the data mover 61 over a high-speed point-to-point bypass data path 66. In response to a request from the client 64 for access to a data file, the data mover 61 grants a lock to the client 64 and sends to the client 64 metadata of the file including pointers to where the data to be accessed is stored in the file

system 62. The client 64 uses the metadata to send a read/write request to the file system 62 over the existing point-to-point bypass data path 66 (see ¶ 56).

The data mover 61 does not send to the client 64 routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to the request to access the data file. As explained above, Vahalia teaches that (¶ 52):

The term metadata refers to information about the data, and the term metadata is inclusive of file access information and file attributes. The file access information includes the locks upon the files or blocks of data in the files. The file attributes include pointers to where the data is stored in the cached disk array.

Thus, the metadata sent by the data mover 61 does not include routing metadata comprising one or more node addresses along one or more network routes between the client and the server. Moreover, the inclusion of such routing metadata in the transmission from the data mover 61 in response to the client's file access request would not serve any purpose whatsoever in the context of the embodiment of FIG. 3 because physical addresses and IP addresses are ignored in communications over point-to-point, high-speed data links in accordance with the knowledge that was generally available at the time the invention was made.

e. The embodiment disclosed in FIG. 4 of Vahalia

The embodiment shown in FIG. 4 combines the features of the embodiments of FIGS. 2 and 3. In particular, a client 88 connects to a first data mover 81 over a network connection, connects to a file system 83 owned by the first data mover 81 over a high-speed point-to-point bypass data path 92, and connects to a file system 84 owned by a second data mover 82 over a high-speed point-to-point bypass data path 93. The client 88 obtains access to a data file stored in file system 83 over the point-to-point bypass data path 92 in the manner described above in connection with the embodiment of FIG. 3. The client 88 obtains access to a data file stored in the file system 84 over the point-to-point bypass data path 93 in an analogous way, except that the metadata request from the client 88 is forwarded from the first data mover 81 to the second data mover 82 (see ¶ 62).

The data mover 81 does not send to the client 88 routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to the request to access the data file. As explained above, Vahalia teaches that (¶ 52):

The term metadata refers to information about the data, and the term metadata is inclusive of file access information and file attributes. The file access information includes the locks upon the files or blocks of data in the files. The file attributes include pointers to where the data is stored in the cached disk array.

Thus, the metadata sent by the data mover 81 does not include routing metadata comprising one or more node addresses along one or more network routes between the client and the server. Moreover, the inclusion of such routing metadata in the transmission from the data mover 81 in response to the client's file access request would not serve any purpose whatsoever in the context of the embodiment of FIG. 4 because physical addresses and IP addresses are ignored in communications over point-to-point, high-speed data links 92, 93 in accordance with the knowledge that was generally available at the time the invention was made.

e. Conclusion

In summary, Vahalia does not disclose a method that includes sending from the source site to a client site routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site in response to a request from the client site for access to a data file stored in one or more physical storage systems at the source site, as recited in claim 1.

For at least these reasons, the Examiner's rejection of independent claim 1 under 35 U.S.C. § 102(e) over Vahalia now should be withdrawn.

2. Dependent claims 2 and 6

Each of claims 2 and 6 incorporates the features of independent claim 1 and therefore is patentable over Vahalia for at least the same reasons explained above.

3. Claims 12, 13, and 19

The pertinent features of independent claims 12 and 19 essentially track the features of independent claim 1 discussed above. Therefore, claims 12 and 19 are patentable over Vahalia for at least the same reasons explained above.

Claim 13 incorporates the features of independent claim 12 and therefore is patentable over Vahalia for at least the same reasons.

B. Rejection under 35 U.S.C. § 103(a) over Vahalia in view of Koyanagi (U.S. 20010013067)

The Examiner has rejected claims 4, 5, 23, 24, 26, and 27 under 35 U.S.C. § 103(a) over Vahalia in view of Koyanagi (U.S. 20010013067).

Claims 4 and 5 incorporate the features of independent claim 1 and claims 23, 24, 26, and 27 incorporate the features of independent claim 12. Koyanagi does not make-up for the failure of Vahalia to teach the features of independent claim 1 discussed above. Indeed, the Examiner has cited Koyanagi merely for the proposition that routing meta data comprises a next hop address (see page 5 of the final Office action). Therefore, claims 4, 5, 23, 24, 26, and 27 are patentable over Vahalia in view of Koyanagi for at least the same reasons explained above in connection with independent claim 1.

C. Rejection under 35 U.S.C. § 102(b) over Vahalia in view of Kato (U.S. 6,223,249)

The Examiner has rejected claims 21, 22, and 25 under 35 U.S.C. § 103(a) over Vahalia in view of Kato (U.S. 6,223,249).

Claim 21 incorporates the features of independent claim 1, claim 22 incorporates the features of independent claim 12, and claim 25 incorporates the features of independent claim 19. Kato does not make-up for the failure of Vahalia to teach the features of independent claim 1 discussed above. Indeed, the Examiner has cited Kato merely for showing in FIGS. 10A and 10B a block map that arranges data by disk number and sector number (see page 6 of the final Office action). Therefore, claims 21, 22, and 25 are patentable over Vahalia in view of Kato for at least the same reasons explained above in connection with claim 1.

Applicant : Lance W. Russell
Serial No. : 09/888,544
Filed : June 25, 2001
Page : 15 of 20

Attorney's Docket No.: 10003533-1
Appeal Brief dated May 5, 2006
Reply to final action dated Dec. 2, 2005

VIII. Conclusion

For the reasons explained above, all of the pending claims are now in condition for allowance and should be allowed.

Charge any excess fees or apply any credits to Deposit Account No. 08-2025.

Respectfully submitted,

Date: May 5, 2006



Edouard Garcia
Reg. No. 38,461
Telephone No.: (650) 289-0904

Please direct all correspondence to:

Hewlett-Packard Company
Intellectual Property Administration
Legal Department, M/S 35
P.O. Box 272400
Fort Collins, CO 80528-9599

CLAIMS APPENDIX

The claims that are the subject of Appeal are presented below.

Claim 1 (previously presented): A method of accessing a data file in a distributed computing environment, comprising:

in response to a request from a client site for access to a data file stored in one or more physical storage systems at a source site, sending from the source site to the client site physical address meta data including physical addresses of one or more logical blocks of the data file in the one or more physical storage systems, and routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site.

Claim 2 (previously presented): The method of claim 1, further comprising storing at the source site a data structure comprising the physical address meta data and the routing meta data for one or more logical file blocks of the requested data file.

Claim 3 (canceled)

Claim 4 (previously presented): The method of claim 1, wherein the routing meta data comprises next hop node addresses from the client site for each of the one or more network routes.

Claim 5 (previously presented): The method of claim 1, wherein the routing meta data comprises complete path information from the client site to the source site for each of the one or more network routes.

Claim 6 (original): The method of claim 1, wherein the meta data is sent to the client site in accordance with a routable network protocol.

Claims 7-11 (canceled)

Claim 12 (previously presented): A system for accessing a data file in a distributed computing environment, comprising:

a file system of a source site configured to manage access to one or more logical file blocks of a data file stored in one or more physical storage systems of the source site, wherein, in response to a request from a client site for access to the data file, the file system sends from the source site to the client site physical address meta data including physical addresses of one or more logical blocks of the data file in the one or more physical storage systems, and routing meta data comprising one or more node addresses along one or more network routes between the client site and the source site.

Claim 13 (previously presented): The system of claim 12, wherein the file system is configured to store at the source site a data structure comprising the physical address meta data and the routing meta data for one or more logical file blocks of the requested data file.

Claims 14-18 (canceled)

Claim 19 (previously presented): A machine-readable medium encoded with a data structure for accessing a data file in a distributed computing environment, comprising:

physical address meta data including physical addresses of one or more logical blocks of the data file in one or more physical storage systems of a source site; and
routing meta data comprising one or more node addresses along one or more network routes between a client site and the source site.

Claim 20 (canceled)

Claim 21 (previously presented): The method of claim 1, wherein the physical address meta data comprises physical address parameters including disk number and sector number where one or more logical blocks of the data file are stored in the one or more physical storage systems.

Claim 22 (previously presented): The system of claim 12, wherein the physical address meta data comprises physical address parameters including disk number and sector

number where one or more logical blocks of the data file are stored in the one or more physical storage systems.

Claim 23 (previously presented): The system of claim 12, wherein the routing meta data comprises next hop node addresses from the client site for each of the one or more network routes.

Claim 24 (previously presented): The system of claim 12, wherein the routing meta data comprises complete path information from the client site to the source site for each of the one or more network routes.

Claim 25 (previously presented): The machine-readable medium of claim 19, wherein the physical address meta data comprises physical address parameters including disk number and sector number where one or more logical blocks of the data file are stored in the one or more physical storage systems.

Claim 26 (previously presented): The machine-readable medium of claim 12, wherein the routing meta data comprises next hop node addresses from the client site for each of the one or more network routes.

Claim 27 (previously presented): The machine-readable medium of claim 12, wherein the routing meta data comprises complete path information from the client site to the source site for each of the one or more network routes.

Applicant : Lance W. Russell
Serial No. : 09/888,544
Filed : June 25, 2001
Page : 19 of 20

Attorney's Docket No.: 10003533-1
Appeal Brief dated May 5, 2006
Reply to final action dated Dec. 2, 2005

EVIDENCE APPENDIX

There is no evidence submitted pursuant to 37 CFR §§ 1.130, 1.131, or 1.132 or any other evidence entered by the Examiner and relied upon by Appellant in the pending appeal. Therefore, no copies are required under 37 CFR § 41.37(c)(1)(ix) in the pending appeal.

Applicant : Lance W. Russell
Serial No. : 09/888,544
Filed : June 25, 2001
Page : 20 of 20

Attorney's Docket No.: 10003533-1
Appeal Brief dated May 5, 2006
Reply to final action dated Dec. 2, 2005

RELATED PROCEEDINGS APPENDIX

Appellant is not aware of any decisions rendered by a court or the Board that will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal. Therefore, no copies are required under 37 CFR § 41.37(c)(1)(x) in the pending appeal.